# Adjoint Methods 

Steven G. Johnson

Created Spring 2006, updated December 17, 2012.

## 1 Introduction

Given the solution $\mathbf{x}$ of a discretized PDE or some other set of $M$ equations parameterized by $P$ variables $\mathbf{p}$ (design parameters, a.k.a. control variables or decision parameters), we often wish to compute some function $g(\mathbf{x}, \mathbf{p})$ based on the parameters and the solution. For example, if the PDE is a wave equation, we might want to know the scattered power in some direction. Or, for a mechanical simulation, we might want to know the load-bearing capacity of the structure. Or for a fluid, we might wish to know the flow rate somewhere. Often, however, we want to know more than just the value of $g$-we also want to know its gradient $\frac{d g}{d \mathbf{p}}$. Adjoint methods give an efficient way to evaluate $\frac{d g}{d \mathbf{p}}$, with a cost independent of $P$ and usually comparable to the cost of solving for $\mathbf{x}$ once.

The gradient of $g$ with respect to $\mathbf{p}$ is extremely useful. It gives a measure of the sensitivity of our answer to the parameters $\mathbf{p}$ (which may, for example, come from some experimental measurements with some associated uncertainties). Or, we may want to perform an optimization of $g$, picking the $\mathbf{p}$ that produce some desired result; in this case the gradient indicates a useful search direction (e.g. for nonlinear conjugate-gradient optimization). For largescale optimization problems, the number $P$ of design parameters can be hundreds, thousands, or morethis is common in shape or topology optimization, in which $\mathbf{p}$ controls the placement and shape of arbitrary blobs of different materials constituting a given structure/design. Sometimes, this process is called inverse design: finding the problem that yields a given solution instead of the other way around. When $P \gg 1$, the amazing efficiency of adjoint methods makes inverse design possible.

I hadn't found any textbook description of adjoint methods that I particularly like, which is part of my motivation for writing up these notes. One introduction can be found in [1], and a more general treatment
can be found in [2]. Subsequently, Gil Strang wrote a nice introduction to adjoint methods in his book [3], including a discussion of the important topic of automatic differentiation (for which adjoint or "reverse" differentiation is a key idea).

## 2 Linear equations

Suppose that the column-vector $\mathbf{x}$ solves the $M \times M$ linear equation $A \mathbf{x}=\mathbf{b}$ where we take $\mathbf{b}$ and $A$ to be real ${ }^{1}$ and to depend in some way on $\mathbf{p}$. To evaluate the gradient directly, we would do

$$
\frac{d g}{d \mathbf{p}}=g_{\mathbf{p}}+g_{\mathbf{x}} \mathbf{x}_{\mathbf{p}}
$$

where the subscripts indicate partial derivatives ( $g_{\mathbf{x}}$ is a row vector, $\mathbf{x}_{\mathbf{p}}$ is an $M \times P$ matrix, etc.). Since $g$ is a given function, $g_{\mathbf{p}}$ and $g_{\mathbf{x}}$ are presumably easy to compute. On the other hand, computing $\mathbf{x}_{\mathbf{p}}$ is hard: evaluating it directly by differentiating $A \mathbf{x}=\mathbf{b}$ by a parameter $p_{i}$ gives $\mathbf{x}_{p_{i}}=A^{-1}\left(\mathbf{b}_{p_{i}}-A_{p_{i}} \mathbf{x}\right)$. That is, we would have to solve an $M \times M$ linear equation for $P$ right-hand sides, once for every compont of $\mathbf{p}$; this is impractical if $P$ and $M$ are large.

More explicitly, the problematic term is:
$g_{\mathbf{x}} \mathbf{x}_{\mathbf{p}}=\underbrace{g_{\mathbf{x}}}_{1 \times M}[\underbrace{A^{-1}}_{M \times M} \underbrace{\left(\mathbf{b}_{\mathbf{p}}-A_{\mathbf{p}} \mathbf{x}\right)}_{M \times P}]=\underbrace{\left[g_{\mathbf{x}} A^{-1}\right]}_{1 \times M} \underbrace{\left(\mathbf{b}_{\mathbf{p}}-A_{\mathbf{p}} \mathbf{x}\right)}_{M \times P}$,
where $A_{\mathbf{p}} \mathbf{x}$ denotes the $M \times P$ matrix with columns $A_{p_{i}} \mathbf{x}$ for $i=1, \ldots P .^{2}$ One way of looking at the difficulty is that in the first equation we multiply a $M \times M$ matrix by a $M \times P$ matrix, which costs $O\left(M^{2} P\right)$ work, or equivalently we have multiplications of $A^{-1}$

[^0]by $P$ vectors (i.e., solves of $P$ right-hand sides, which in practice would likely use a factorization of $A$ or an iterative solver rather than explicitly computing $\left.A^{-1}\right) .^{3}$ However, this can be ameliorated simply by parenthesizing in a different way [3], ${ }^{4}$ as shown in the last expression. If we multiply $\lambda^{T}=g_{\mathbf{x}} A^{-1}$ first, that corresponds to only a single solution of an adjoint equation ${ }^{5}$
\[

$$
\begin{equation*}
A^{T} \boldsymbol{\lambda}=g_{\mathbf{x}}^{T} \tag{1}
\end{equation*}
$$

\]

and then we multiply a single vector $\lambda^{T}$ by our $M \times P$ matrix for only $\theta(M P)$ work. Putting it all together, we obtain:

$$
\frac{d g}{d \mathbf{p}}{ }_{\mathbf{f}=0}=g_{\mathbf{p}}-\boldsymbol{\lambda}^{T} \mathbf{f}_{\mathbf{p}}=g_{\mathbf{p}}-\boldsymbol{\lambda}^{T}\left(A_{\mathbf{p}} \mathbf{x}-\mathbf{b}_{\mathbf{p}}\right) .
$$

Again, $A(\mathbf{p})$ and $\mathbf{b}(\mathbf{p})$ are presumably specified analytically and thus $A_{\mathbf{p}}$ and $\mathbf{b}_{\mathbf{p}}$ can easily be computed (in some cases automatically, by automatic program differentiators such as ADIFOR). Note that the adjoint problem is of the same size as the original $A \mathbf{x}=\mathbf{b}$ system, can use the same factorization (e.g. LU factorization $A=L U$ immediately gives $A^{T}=U^{T} L^{T}$ ), has the same condition number, and has the same spectrum of eigenvalues (the eigenvalues of $A$ and $A^{T}$ are identical) so iterative algorithms will have similar performance (and can use similar preconditioners)—in every sense, solving the adjoint problem should be no harder than solving the original problem.

## 3 Nonlinear equations

If $\mathbf{x}$ satisfies some general, possibly nonlinear, equations $\mathbf{f}(\mathbf{x}, \mathbf{p})=0$, the process is almost exactly the same. Differentiating the $\mathbf{f}$ equation, we find $\mathbf{f}_{\mathbf{x}} \mathbf{x}_{\mathbf{p}}+$

[^1]$\mathbf{f}_{\mathbf{p}}=0$ and thus $\mathbf{x}_{p}=-\mathbf{f}_{\mathbf{x}}^{-1} \mathbf{f}_{\mathbf{p}}$. Hence, we write
$\frac{d g}{d \mathbf{p}}=g_{\mathbf{p}}+g_{\mathbf{x}} \mathbf{x}_{\mathbf{p}}=g_{\mathbf{p}}-\underbrace{g_{\mathbf{x}}}_{1 \times M}[\underbrace{\mathbf{f}_{\mathbf{x}}^{-1}}_{M \times M} \underbrace{\mathbf{f}_{\mathbf{p}}}_{M \times P}]=g_{\mathbf{p}}-\underbrace{\left[g_{\mathbf{x}} \mathbf{f}_{\mathbf{x}}^{-1}\right]}_{1 \times M} \underbrace{\mathbf{f}_{\mathbf{p}}}_{M \times P}$.
We solve for $\mathbf{x}$ by whatever method, then solve for $\boldsymbol{\lambda}$ from
\[

$$
\begin{equation*}
\mathbf{f}_{\mathbf{x}}^{T} \boldsymbol{\lambda}=\mathbf{g}_{\mathbf{x}}^{T} \tag{2}
\end{equation*}
$$

\]

and finally obtain

$$
\begin{equation*}
\frac{d g}{d \mathbf{p}}_{\mathbf{f}=0}=g_{\mathbf{p}}-\boldsymbol{\lambda}^{T} \mathbf{f}_{\mathbf{p}} \tag{3}
\end{equation*}
$$

The only difference is that the adjoint equation (2) is not simply the adjoint of the equation for $\mathbf{x}$. Still, it is a single $M \times M$ linear equation for $\boldsymbol{\lambda}$ that should be of comparable (or lesser) difficulty to solving for $\mathbf{x}$.

## 4 Eigenproblems

As a more complicated example illustrating the use of equations (2) and (3) from the previous sections, let us suppose that we are solving a linear eigenproblem $A \mathbf{x}=\alpha \mathbf{x}$ and looking at some function $g(\mathbf{x}, \alpha, \mathbf{p})$. For simplicity, assume that $A$ is realsymmetric and that $\alpha$ is simple (non-degenerate; i.e., $\mathbf{x}$ is the only eigenvector for $\alpha$ ). ${ }^{6}$ In this case, we now have $M+1$ unknowns described by the column vector:

$$
\tilde{\mathbf{x}}=\binom{\mathbf{x}}{\alpha}
$$

The eigenequation $\mathbf{f}=A \mathbf{x}-\alpha \mathbf{x}$ only gives us $M$ equations and doesn't completely determine $\tilde{\mathbf{x}}$, for two reasons. First, of course, there are many possible eigenvalues, but let's assume that we have picked one in some fashion (e.g. the smallest $\alpha$, or the $\alpha$ closest to $\pi$, or the third largest $|\alpha|$, or ...). Second, the eigenequation does not determine the length $|\mathbf{x}|$; let's arbitrarily pick $|\mathbf{x}|=1$ or $\mathbf{x}^{T} \mathbf{x}=1$. This gives us $M+1$ equations $\tilde{\mathbf{f}}=0$ where:

$$
\tilde{\mathbf{f}}=\binom{\mathbf{f}}{\mathbf{x}^{T} \mathbf{x}-1} .
$$

[^2]We'll need $M+1$ adjoint variables $\tilde{\lambda}$ :

$$
\tilde{\lambda}=\binom{\lambda}{\beta}
$$

The adjoint equations (2) then give:

$$
\begin{gather*}
(A-\alpha) \boldsymbol{\lambda}=g_{\mathbf{x}}^{T}-2 \beta \mathbf{x}  \tag{4}\\
-\mathbf{x}^{T} \boldsymbol{\lambda}=g_{\alpha} \tag{5}
\end{gather*}
$$

The first equation, at first glance, seems to be problematic: $A-\alpha$ is singular, with a null space of $\mathbf{x}$. It's, okay, though! First, we have to choose $\beta$ so that solutions of equation (4) exist: the right-hand side must be orthogonal to $\mathbf{x}$ so that it is not in the null space of $A-\alpha$. That is, we must have $\mathbf{x}^{T}\left(g_{\mathbf{x}}^{T}-2 \beta \mathbf{x}\right)=0$, and thus $\beta=\mathbf{x}^{T} g_{\mathbf{x}}^{T} / 2$ (since $\mathbf{x}^{T} \mathbf{x}=1$ ), and therefore $\boldsymbol{\lambda}$ satisfies:

$$
\begin{equation*}
(A-\alpha) \boldsymbol{\lambda}=\left(1-\mathbf{x x}^{T}\right) g_{\mathbf{x}}^{T}=P g_{\mathbf{x}}^{T} \tag{6}
\end{equation*}
$$

where $P=1-\mathbf{x} \mathbf{x}^{T}$ is the projection operator into the space orthogonal to $\mathbf{x}$. This equation then has a solution, and in fact it has infinitely many solutions: we can add any multiple of $\mathbf{x}$ to $\boldsymbol{\lambda}$ and still have a solution. Equivalently, we can write $\boldsymbol{\lambda}=\boldsymbol{\lambda}_{0}+\gamma \mathbf{x}$ for $\mathbf{x}^{T} \boldsymbol{\lambda}_{0}=0$ and some $\gamma$. Fortunately, $\gamma$ is determined by (5): $\gamma=-g_{\alpha}$. Finally, with $\boldsymbol{\lambda}_{0}$ determined by (6), ${ }^{7}$ we can find the desired gradient via (3):

$$
\begin{equation*}
\frac{d g}{d \mathbf{p}}_{\mathbf{f}=0}=g_{\mathbf{p}}-\boldsymbol{\lambda}^{T} A_{p} \mathbf{x}=g_{\mathbf{p}}-\boldsymbol{\lambda}_{0}^{T} A_{p} \mathbf{x}+g_{\alpha} \mathbf{x}^{T} A_{p} \mathbf{x} \tag{7}
\end{equation*}
$$

If we compare with $\frac{d g}{d \mathbf{p}}=g_{\mathbf{p}}+g_{\mathbf{x}} \mathbf{x}_{\mathbf{p}}+g_{\alpha} \alpha_{\mathbf{p}}$, we immediately see that $\alpha_{\mathbf{p}}=\mathbf{x}^{T} A_{p} \mathbf{x}$. This is a well-known result from quantum physics and perturbation theory, where it is known as the Hellman-Feynman theorem.

## 5 Example inverse design

As a more concrete example of an inverse-design problem, let's consider the Schrodinger eigenequation in one dimension,

$$
\left[-\frac{d^{2}}{d x^{2}}+V(x)\right] \psi(x)=E \psi(x)
$$

with periodic boundaries $\psi(x+2)=\psi(x)$. Normally, we take a given $V(x)$ and solve for $\psi$ and $E$.

[^3]

Figure 1: Optimized $V(x)$ (scaled by $1 / 1000$ ) and $\psi(x)$ for $\psi_{0}(x)=1+\sin [\pi x+\cos (3 \pi x)]$ after 500 cg iterations.

Now, however, we will specify a particular $\psi_{0}(x)$ and find the $V(x)$ that gives $\psi(x) \approx \psi_{0}(x)$ for the groundstate eigenfunction (i.e. for the smallest eigenvalue $E)$. In particular, we will find the $V(x)$ that minimizes

$$
g=\int_{-1}^{1}\left|\psi(x)-\psi_{0}(x)\right|^{2} d x
$$

To solve this numerically, we will discretize the interval $x \in[-1,1)$ with $M$ equally-spaced points $x_{n}=$ $n \Delta x\left(\Delta x=\frac{2}{M+1}\right)$, and solve for the solution $\psi\left(x_{n}\right)$ at these points, denoted by the vector $\boldsymbol{\psi}$. That is, to compare with the notation of the previous sections, we have the eigenvector $\mathbf{x}=\boldsymbol{\psi}$, the eigenvalue $\alpha=E$, and the parameters $V\left(x_{n}\right)$ or $\mathbf{p}=\mathbf{V}$. If we discretize the eigenoperator with the usual centerdifference scheme, we get $A \boldsymbol{\psi}=E \boldsymbol{\psi}$ for:

$$
A=\frac{1}{\Delta x^{2}}\left(\begin{array}{cccccc}
2 & -1 & 0 & \cdots & 0 & -1 \\
-1 & 2 & -1 & 0 & \cdots & \\
0 & -1 & 2 & -1 & 0 & \cdots \\
\vdots & & & \ddots & & \\
& & & -1 & 2 & -1 \\
-1 & 0 & \cdots & 0 & -1 & 2
\end{array}\right)+\operatorname{diag}(\mathbf{V}) .
$$

As before, we normalize $\boldsymbol{\psi}$ (and $\boldsymbol{\psi}_{0}$ ) to $\boldsymbol{\psi}^{T} \boldsymbol{\psi}=1,{ }^{8}$ giving a projection operator $P=1-\boldsymbol{\psi} \boldsymbol{\psi}^{T}$ (or $P=$ $1-|\psi\rangle\langle\psi|$, in Dirac notation). The discrete version of $g$ is now

$$
g(\boldsymbol{\psi}, \mathbf{V})=\left(\boldsymbol{\psi}-\boldsymbol{\psi}_{0}\right)^{T}\left(\boldsymbol{\psi}-\boldsymbol{\psi}_{0}\right) \Delta x
$$

[^4]where $\boldsymbol{\psi}_{0}$ is $\psi_{0}\left(x_{n}\right)$, our target eigenfunction. Therefore, $g_{\boldsymbol{\psi}}=2\left(\boldsymbol{\psi}-\boldsymbol{\psi}_{0}\right)^{T} \Delta x$ and thus, by eq. (6), we find $\boldsymbol{\lambda}$ via:
\[

$$
\begin{equation*}
(A-E) \boldsymbol{\lambda}=2 P\left(\boldsymbol{\psi}-\boldsymbol{\psi}_{0}\right) \Delta x \tag{8}
\end{equation*}
$$

\]

with $P \boldsymbol{\lambda}=0\left(\boldsymbol{\lambda}=\boldsymbol{\lambda}_{0}\right.$ since $\left.g_{E}=0\right) . g_{\mathbf{V}}$ and $g_{E}$ are both 0 . Moreover, $A_{V_{n}}$ is simply the matrix with 1 at ( $n, n$ ) and 0's elsewhere, and thus from (7):

$$
\frac{d g}{d V_{n}}=-\lambda_{n} \psi_{n}
$$

or equivalently $\frac{d g}{d \boldsymbol{V}}=-\boldsymbol{\lambda} \quad \boldsymbol{\psi}$ where $\quad$ is the pointwise product (. $*$ in Matlab).

Whew! Now how do we solve these equations numerically? This is illustrated by the Matlab function schrodinger_fd_adj given below. We set up $A$ as a sparse matrix, then find the smallest eigenvalue and eigenvector via the eigs function (which uses an iterative Arnoldi method). Then we solve (8) for $\boldsymbol{\lambda}$ via the Matlab pcg function (preconditioned conjugategradient, although we don't bother with a preconditioner).

Then, given $g$ and $\frac{d g}{d \mathbf{V}}$, we then just plug it into some optimization algorithm. In particular, nonlinear conjugate gradient seems to work well for this problem. ${ }^{9}$

### 5.1 Optimization results

In this section, we give a few example results from running the above procedure (nonlinear cg optimization) for $M=100$. As the starting guess for our optimization, we'll just use $V(x)=0$. That is, we are doing a local optimization in a 100-dimensional space, using the adjoint method to get the gradient. It is somewhat remarkable that this works-in a few seconds on a PC, it converges to a very good solution!

We'll try a couple of example $\psi_{0}(x)$ functions. To start with, let's do $\psi_{0}(x)=1+\sin [\pi x+\cos (3 \pi x)]$. (Note that the ground-state $\psi$ will never have any nodes, so we require $\psi_{0} \geq 0$ everywhere.) This $\psi_{0}(x)$, along with the resulting $\psi(x)$ and $V(x)$ after 500 cg iterations, are shown in figure 1 . The solution $\psi(x)$ matches $\psi_{0}(x)$ very well except for a couple of small ripples, and $V(x)$ is quite complicated-not something you could easily guess!

[^5]

Figure 2: Optimized $V(x)$ (scaled by $1 / 10000$ ) and $\psi(x)$ for $\psi_{0}(x)=1-|x|$ for $|x|<0.5$, after 5000 cg iterations.


Figure 3: Optimized $\psi(x)$ for $\psi_{0}(x)=1-|x|$ for $|x|<0.5$, after various numbers of nonlinear conjugate-gradient iterations (from 10 to 10000).

Oh, but that $\psi_{0}$ was too easy! Let's try one with discontinuities: $\psi_{0}(x)=1-|x|$ for $|x|<0.5$ and 0 otherwise (which looks a bit like a "house"). This $\psi_{0}(x)$, along with the resulting $\psi(x)$ and $V(x)$ after 500 cg iterations, are shown in figure 2. Amazingly, it still captures $\psi_{0}$ pretty well, although it has a bit more trouble with the discontinuities than with the slope discontinuity. This time, we let it converge for 5000 cg iterations to give it a bit more time. Was this really necessary? In figure 3 , we plot $\psi(x)$ for 10 , $20,40,80,160,320$, and 5000 cg iterations. It gets the rough shape pretty quickly, but the discontinuous features are converging fairly slowly. (Presumably this could be improved if we found a good preconditioner, or perhaps by a different optimization method or objective function.)

### 5.2 Matlab code

The following code solves for $g$ and $\frac{d g}{d \mathbf{V}}$, not to mention the eigenfunction $\boldsymbol{\psi}$ and the corresponding eigenvalue $E$, for a given $\mathbf{V}$ and $\boldsymbol{\psi}_{0}$.

```
% Usage: [g,gp,E,psi] = schrodinger_fd_adj(x, V, psiO)
%
% Given a column-vector x(:) of N equally spaced x points a
% V of the potential V(x) at those points, solves Schroding
% [ - d^2/dx^2 + V(x) ] psi(x) = E psi(x)
% with periodic boundaries for the lowest "ground state" ei
% wavefunction psi.
%
% Furthermore, it computes the function g = integral |psi -
% the gradient gp = dg/dV (at each point x).
function [g,gp,E,psi] = schrodinger_fd_adj(x, V, psi0)
    dx = x(2) - x(1);
    N = length(x);
    A = spdiags([ones(N,1), -2 * ones(N,1), ones(N,1)], -1:1,
    A(1,N) = 1;
    A(N,1) = 1;
    A = - A/ dx^2 + spdiags(V, O,N,N);
    opts.disp = 0;
    [psi,E] = eigs(A, 1, 'sa', opts);
    E = E(1,1);
    if sum(psi) < 0
        psi = -psi; % pick sign; note that psi' * psi = 1 from
    end
    gpsi = psi - psiO;
    g = gpsi' * gpsi * dx;
    gpsi = gpsi * 2*dx;
    P = @(x) x - psi * (psi' * x); % projection onto directio
    [lambda,flag] = pcg(A - spdiags(E*ones(N,1), O, N,N), P(g
    lambda = P(lambda);
    gp = -real(conj(lambda) .* psi);
    disp(g);
```


## 6 Initial-value problems

So far, we have looked at $\mathbf{x}$ that are determined by "simple" algebraic equations (which may come from a PDE, etcetera). What if, instead, we are determining $\mathbf{x}$ by integrating a set of equations in time? The simplest example of this is an initial-value problem for a linear, time-independent, homogeneous set of ODEs:

$$
\dot{\mathbf{x}}=B \mathbf{x}
$$

whose solution after a time $t$ for $\mathbf{x}(0)=\mathbf{b}$ is formally:

$$
\mathbf{x}=\mathbf{x}(t)=e^{B t} \mathbf{b}
$$

This, however, is exactly a linear equation $A \mathbf{x}=\mathbf{b}$ with $A=e^{-B t}$, so we can just quote our results from earlier! That is, suppose we are optimizing (or evaluating the sensitivity) of some function $g(\mathbf{x}, \mathbf{p})$ based on the solution $\mathbf{x}$ at time $t$. Then we find the adjoint vector $\boldsymbol{\lambda}$ via (1):

$$
e^{-B^{T} t} \boldsymbol{\lambda}=g_{\mathbf{x}}^{T}
$$

Equivalently, $\boldsymbol{\lambda}$ is the exactly the solution $\boldsymbol{\lambda}(t)$ after a time $t$ of its own adjoint ODE:

$$
\dot{\boldsymbol{\lambda}}=B^{T} \boldsymbol{\lambda}
$$

with initial condition $\boldsymbol{\lambda}(0)=g_{\mathbf{x}}^{T}$. We should have expected this by now: solving for $\boldsymbol{\lambda}$ always involves a task of similar complexity to finding $\mathbf{x}$, so if we found $\mathbf{x}$ by integrating an ODE then we find $\boldsymbol{\lambda}$ by an ODE too! Of course, we need not solve these ODEs by matrix exponentials; we can use RungeKutta, forward Euler, or (if $B$ comes from a PDE) whatever scheme we deem appropriate (e.g. CrankNicolson).

One important property to worry about is stability, and here we are in luck. The eigenvalues of $B$ and $B^{T}$ are complex-conjugates, and so if one is stable (eigenvalues with absolute values $\leq 1$ ) then the other is!

Finally, we can write down the gradient $\frac{d g}{d \mathbf{p}}$ via equation (3):

$$
\frac{d g}{d \mathbf{p}}=g_{\mathbf{p}}-\boldsymbol{\lambda}^{T}\left(A_{\mathbf{p}} \mathbf{x}-\mathbf{b}_{\mathbf{p}}\right)
$$

Now, since $A=e^{-B t}$, one might be tempted to write $A_{\mathbf{p}}=-B_{\mathbf{p}} t \cdot A$, but this is not true except in the very special case where $B_{\mathbf{p}}$ commutes with $B$ ! Unfortunately, the general expression for differentiating a
matrix exponential turns out to be more complicated: $A_{\mathbf{p}}=-\int_{0}^{t} e^{-B t^{\prime}} B_{\mathbf{p}} e^{-B\left(t-t^{\prime}\right)} d t^{\prime}$, and so,

$$
\frac{d g}{d \mathbf{p}}=g_{\mathbf{p}}+\int_{0}^{t} \boldsymbol{\lambda}^{T}\left(t-t^{\prime}\right) B_{\mathbf{p}} \mathbf{x}\left(t^{\prime}\right) d t^{\prime}+\boldsymbol{\lambda}^{T} \mathbf{b}_{\mathbf{p}}
$$

This is especially unfortunate because it usually means that we have to store $\mathbf{x}\left(t^{\prime}\right)$ at all times $0 \leq t^{\prime} \leq$ $t$ in order to compute the integral. Adjoint methods are storage-intensive for time-dependent problems!

More generally, of course, one might wish to include time-varying $A$, nonlinearities, inhomogeneous (source) terms, etcetera, into the equations to integrate. A very general formulation of the problem, for differential-algebraic equations (DAEs), can be found in [2]. A similar general principle remains, however: the adjoint variable $\boldsymbol{\lambda}$ is determined by integrating a similar (adjoint) DAE, using the final value of $\mathbf{x}(t)$ to compute the initial condition of $\boldsymbol{\lambda}(0)$. In fact, the $\boldsymbol{\lambda}(t)$ equation is actually often interpreted as being integrated backwards in time from $t$ to 0 . Alternatively, one can consider a "discrete-time" situation of recurrence equations, in which case the adjoint problem is a recurrence "backward in time"see my online notes on adjoint methods for recurrences.

## References

[1] R. M. Errico, "What is an adjoint model?," Bulletin Am. Meteorological Soc., vol. 78, pp. 25772591, 1997.
[2] Y. Cao, S. Li, L. Petzold, and R. Serban, "Adjoint sensitivity analysis for differentialalgebraic equations: The adjoint DAE system and its numerical solution," SIAM J. Sci. Comput., vol. 24, no. 3, pp. 1076-1089, 2003.
[3] G. Strang, Computational Science and Engineering. Wellesley, MA: Wellesley-Cambridge Press, 2007.
[4] A. P. Seyranian, E. Lund, and N. Olhoff, "Multiple eigenvalues in structural optimization problems," Structural Optimization, vol. 8, pp. 207227, 1994.

MIT OpenCourseWare
https://ocw.mit.edu

### 18.335J Introduction to Numerical Methods

Spring 2019

For information about citing these materials or our Terms of Use, visit: https://ocw.mit.edu/terms.


[^0]:    ${ }^{1}$ This involves no loss of generality, since complex linear equations can always be written as real linear equations of twice the size by taking the real and imaginary parts as separate variables.
    ${ }^{2}$ Technically, $A_{\mathbf{p}}$ is a rank-3 tensor or "three-dimensional matrix," although it almost certainly isn't stored this way. For example, $A_{p_{i}} \mathbf{x}$ could be computed for each $i$ separately without saving $A_{p_{i}}$. Often, $A_{p_{i}}$ will be very sparse.

[^1]:    ${ }^{3}$ If $M$ is sparse, then the cost might be significantly less than this $O\left(M^{2} P\right)$ upper bound, but in any case solving $P$ right-hand sides will be significantly more costly than solving a single righthand side for the adjoint formulation.
    ${ }^{4}$ Another way of looking at this, and the source of the $\lambda$ notation, is to think of sort of a "Lagrange multiplier" process: replace $g$ with $\tilde{g}=g-\boldsymbol{\lambda}^{T} \mathbf{f}$ by adding a multiple $\boldsymbol{\lambda}$ of $\mathbf{f}=0$, and then choose $\boldsymbol{\lambda}$ is a clever way to cancel the annoying derivative term. This gives the same result, and may be easier to generalize to some more complicated circumstances, however, such as differential-algebraic equations [2].
    ${ }^{5}$ For complex-valued $\mathbf{x}$ and $A$ and real $g$, instead of the transpose $A^{T}$ one typically obtains the adjoint $A^{\dagger}=A^{T *}$ (the conjugatetranspose).

[^2]:    ${ }^{6}$ Problems involving degenerate eigenvalues occur surprisingly often in optimization of eigenvalues (e.g. when maximizing the minimum eigenvalue of some system), and must be treated with special care. In that case, a generalization of the gradient is required to determine sensitivities or the steepest-descent direction [4], a more elaborate version of what is called degenerate perturbation theory in quantum mechanics [?].

[^3]:    ${ }^{7}$ Since $P$ commutes with $A-\alpha$, we can solve for $\lambda_{0}$ easily by an iterative method such as conjugate gradient: if we start with an initial guess orthogonal to $\mathbf{x}$, all subsequent iterates will also be orthogonal to $\mathbf{x}$ and will thus converge to $\boldsymbol{\lambda}_{0}$ (except for roundoff, which can be corrected by multiplying the final result by $P$ ).

[^4]:    ${ }^{8}$ We also have an arbitrary choice of sign, which we fix by choosing $\int \psi d x>0$.

[^5]:    ${ }^{9}$ I used the nonlinear conjugate-gradient Matlab conj_grad routine from:
    http://www2.imm.dtu.dk/~hbn/Software/

